Steering the Edgecast CDN

Marcel Flores

13 June 2018

verizon

digital media services

The Edgecast CDN

What does it do?

- **The CDN moves content closer to end users.**
- **Reduces latency, increases capacity.**

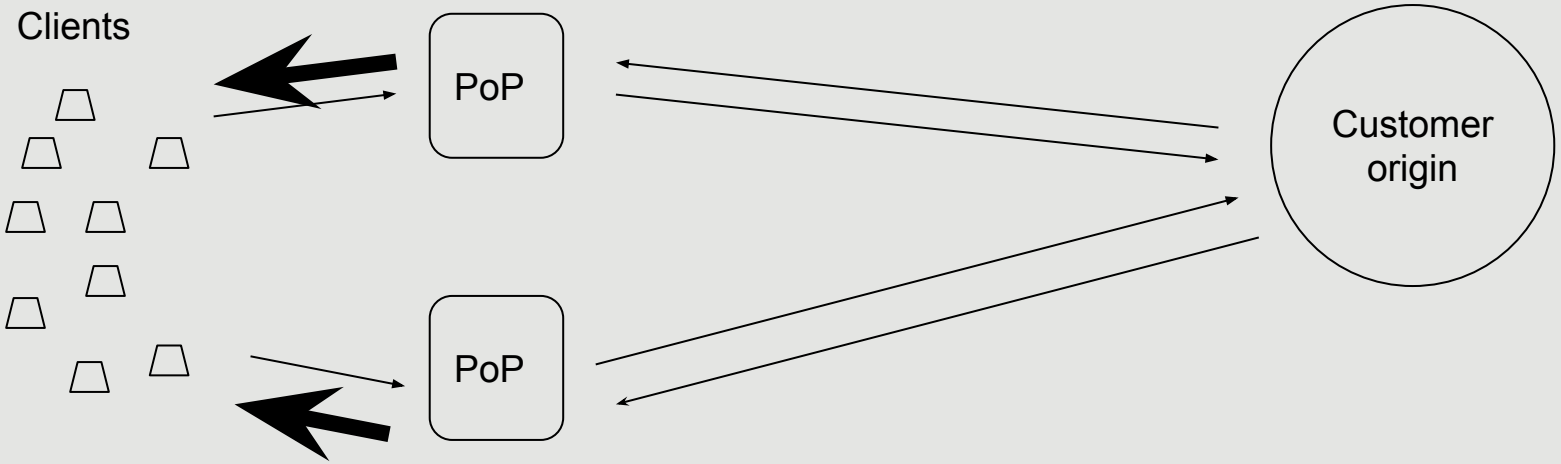

How does it work?

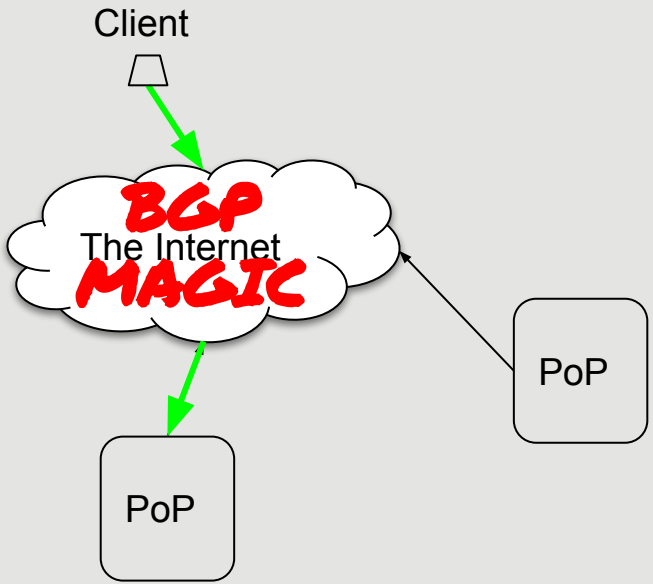

- **Uses anycast for PoP selection.**
- **Relies on BGP to get to the right PoP.**
- **We get lots for free:**
  - Network spreads load more-or-less automatically.
  - We can achieve failover by retracting announcements.

Anycast Challenges

- **Overall, it can be unpredictable:**
  - No information on latency or load.
- **What's going to happen if we change announcements?**
- **Can make many traditional Traffic Engineering/Management problems hard.**

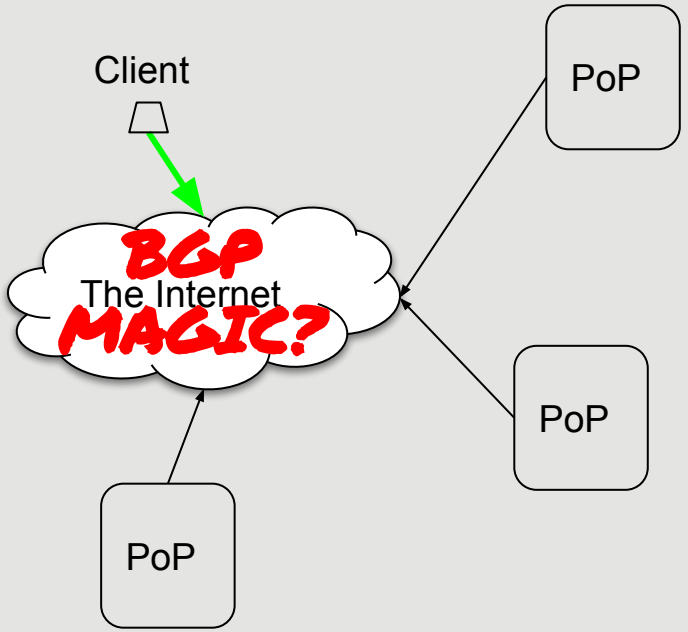

Pulling blocks can be destructive

- **Breaks TCP connections**
  - Unpredictable behavior for a period of time
- **Long running downloads may get ruined**

- **But it can be even worse...**

Pulling blocks can be destructive

A smoother way...

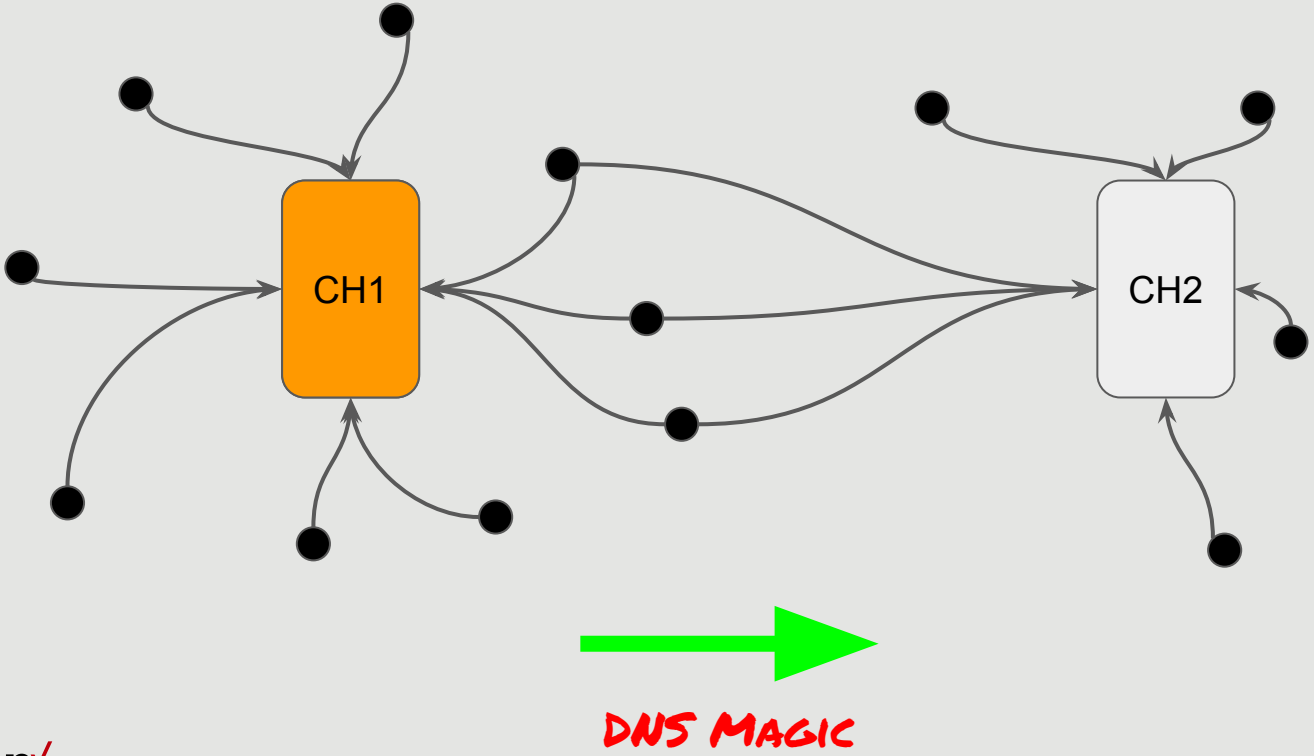

DNS To Different Anycast Announcements

- **Edgecast has a few *networks*.**
- **Each has a (potentially) overlapping set of servers that it addresses.**
  - Kind of like Microsoft's FastRoute
- **Using DNS can steer clients to a particular network.**

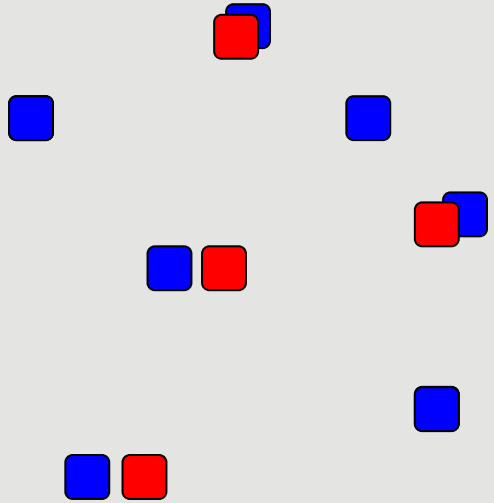

Working with Humans

In the old days...

- **An experienced Human would look at signals:**
  - Current load at PoPs
  - Available capacity at PoPs
- **Human would write and deploy a DNS rule to effect this change.**

In the old days...

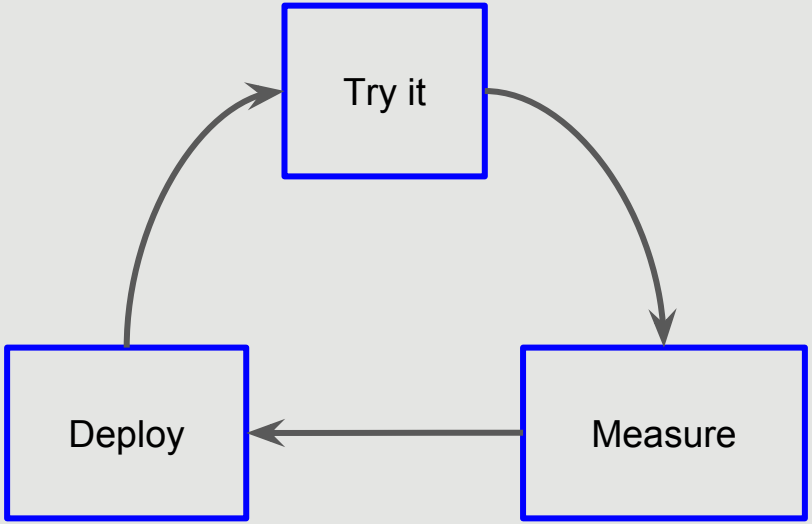

Challenges with humans...

- **Humans make mistakes:**
  - It’s hard to look at a lot of numbers at once,
  - As the CDN grows this gets more serious.
- **Humans have to sleep sometimes!**

How do we make a robot do it?

Heteractis

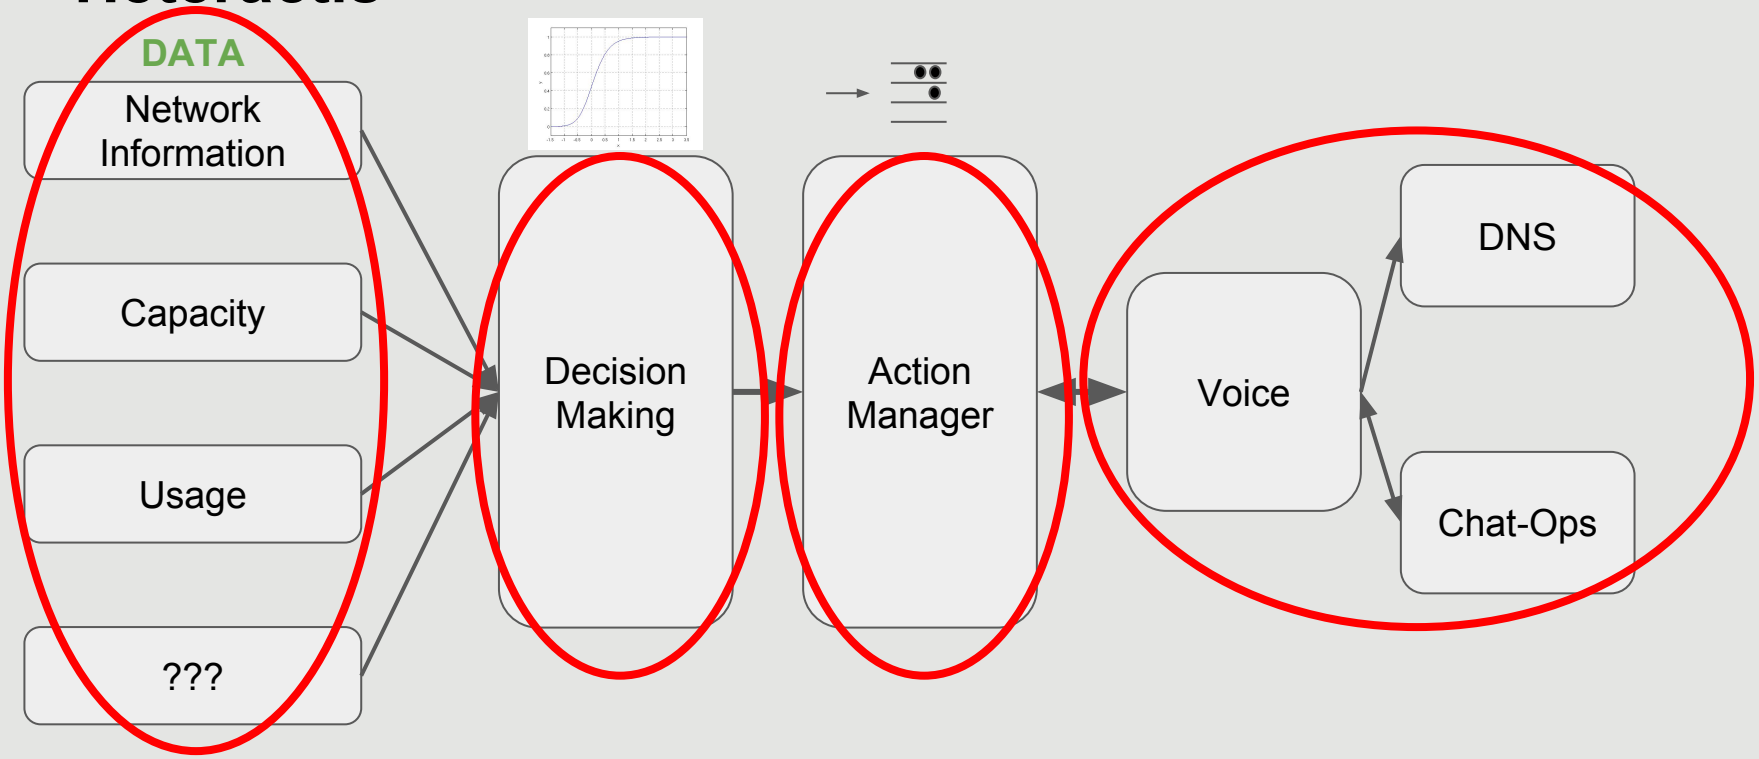
DATA
Network Information
Capacity
Usage
???
Decision Making
Action Manager
Voice
DNS
Chat-Ops

Wow that's complicated.

Heteractis

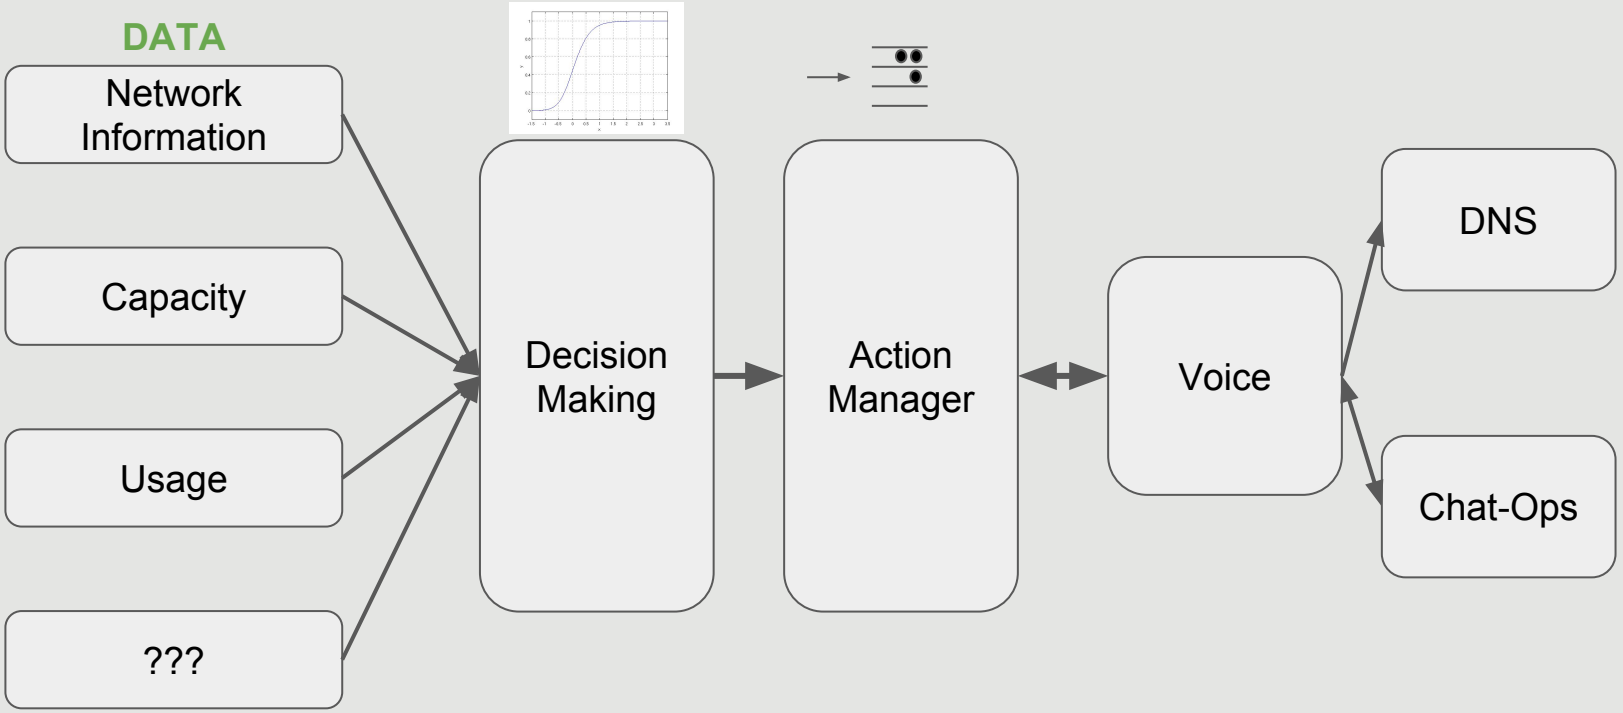
DATA
Network Information
Capacity
Usage
???
Decision Making
Action Manager
Voice
DNS
Chat-Ops

Data Models

- **Collect Many source of data.**
  - Combine into meaningful representations.
  - Capacity, usage, the current state.
- **Keep those up to date.**

Decision Making Models

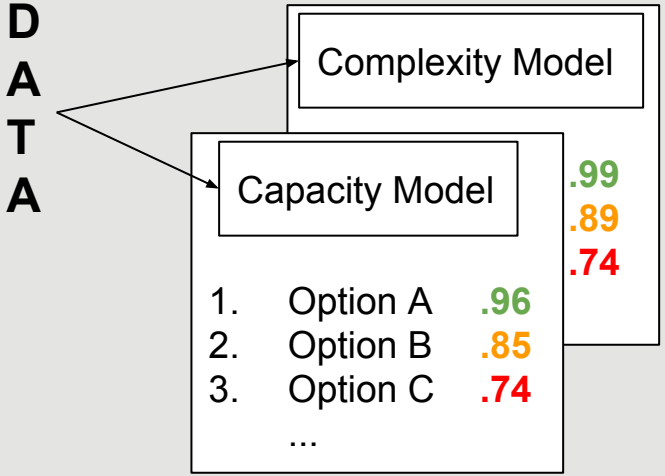

- **Considers a set of possible actions (i.e. DNS rules).**
  - *"For traffic A, send Z% to Red Network"*
- **Asks: According to each model, what would happen if I did this?**
  - Each model generates a score.

Decision Making

- **For each action...**
  - Compute a weighted linear combination.
- **Rank all the actions by combined score.**
- **Pick the action with the highest score.**

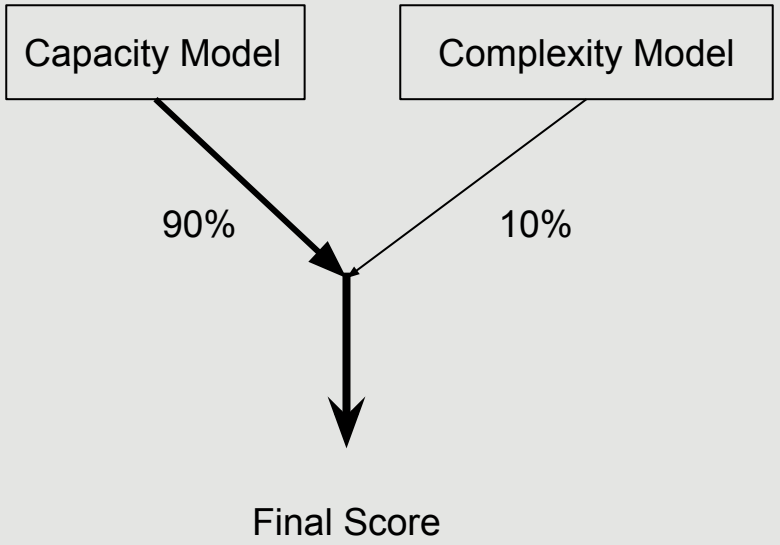

Action Manager

- **Applies some of that delicate touch that a human would do:**
  - Smooths actions out over several minutes.
  - Prevents overlapping changes from firing at once.

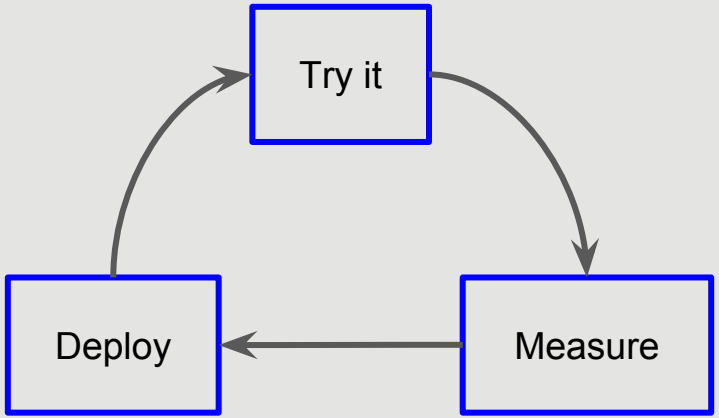

Voice

- **Interacts with humans:**
  - Human gating mechanism: As a deployment strategy, a human says what is OK.
- **Integrates with other systems**
  - Slack, chatops, etc.

Making sure Humans Can understand.

- **Why did it do that?**
  - Can a human validate that it was a good idea?
- **Can ask it questions:**
  - What does Heteractis think about X?
  - Why did it just recommend Y?

Why didn’t we get really fancy?

- **Humans need to feel good about why it’s making decisions.**
- **Each decision:**
  - Can be made based on current data alone.
    - Easy debugging
  - Can be recreated based on snapshot.

Score Models: Capacity

Basic Idea :

- For some proposed action, what would happen?

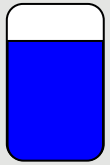

X

- Low Utilization: High Score
- High Utilization: Low Score

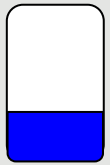

✔

Predicting Load

- **Because we are using anycast, not totally clear what will happen.**
- **So how do we know what score to give it?**

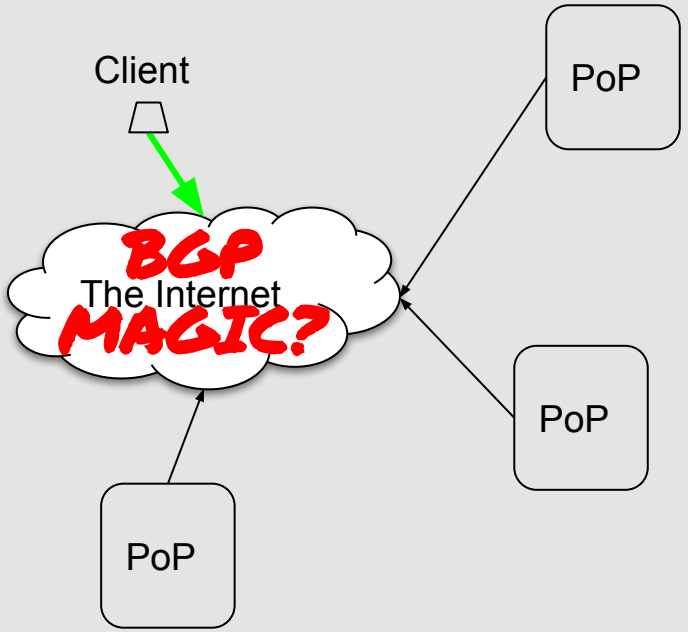

Predicting Load

- **Think about it like fluid:**
  - Sum total amount to move.
  - Distribute it evenly over destinations.

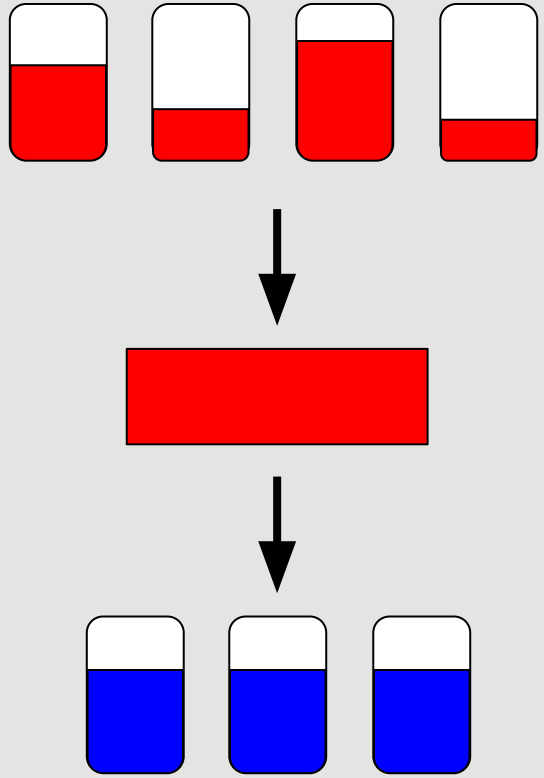

Predicting Load

On the other hand:

- **If we know we already have *X%* at Blue, just scale it!**

At the end: we have an estimate % utilization at each PoP.

Computing Scores

- Use a logistic curve to smooth out the edges.
- Take a harmonic mean of all PoPs to test for outliers below.

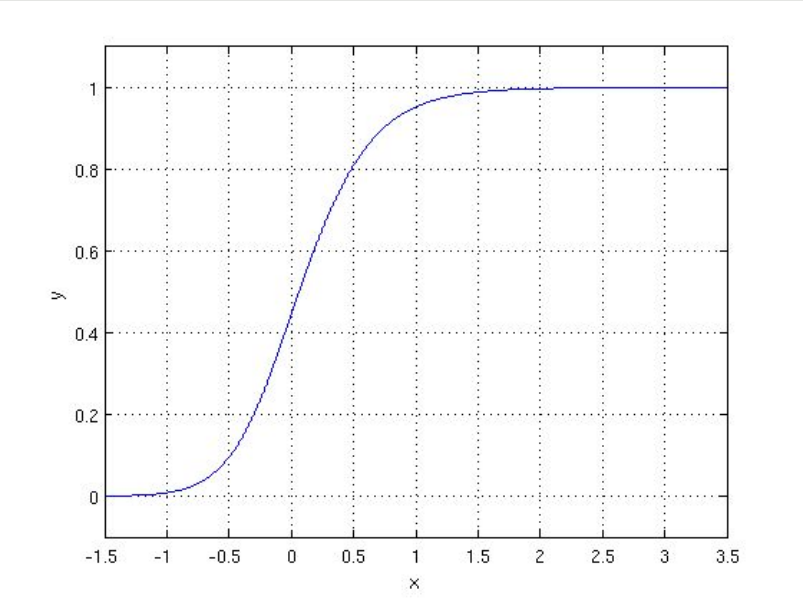

Heteractis In Action

Customer X

Customer X + Daily Traffic

Heteractis's Changes

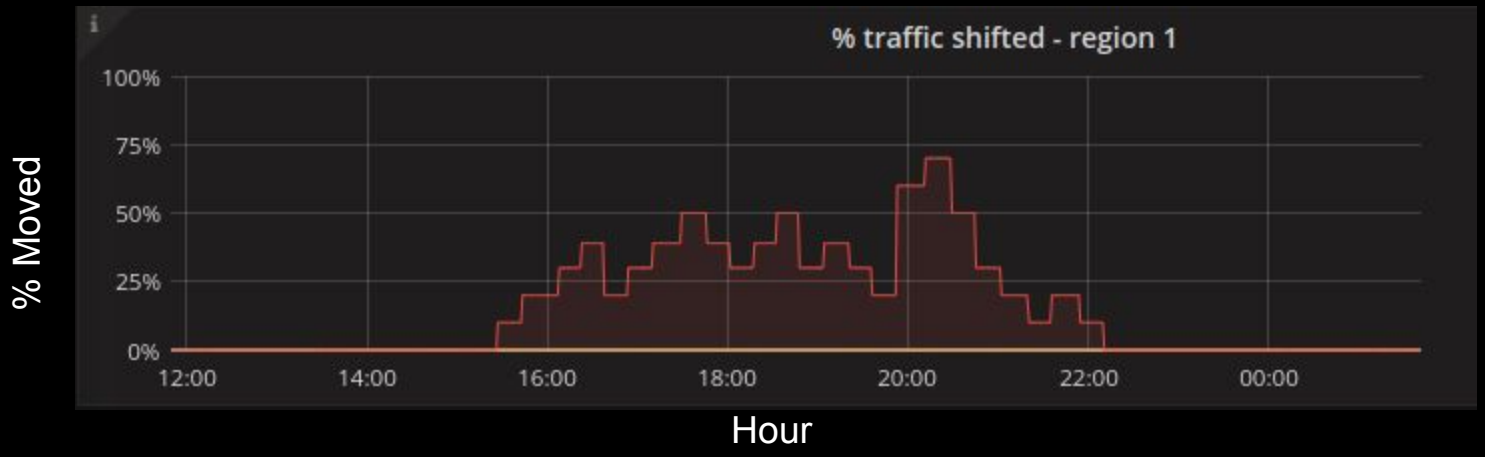

Red Network PoPs

Blue Network PoPs

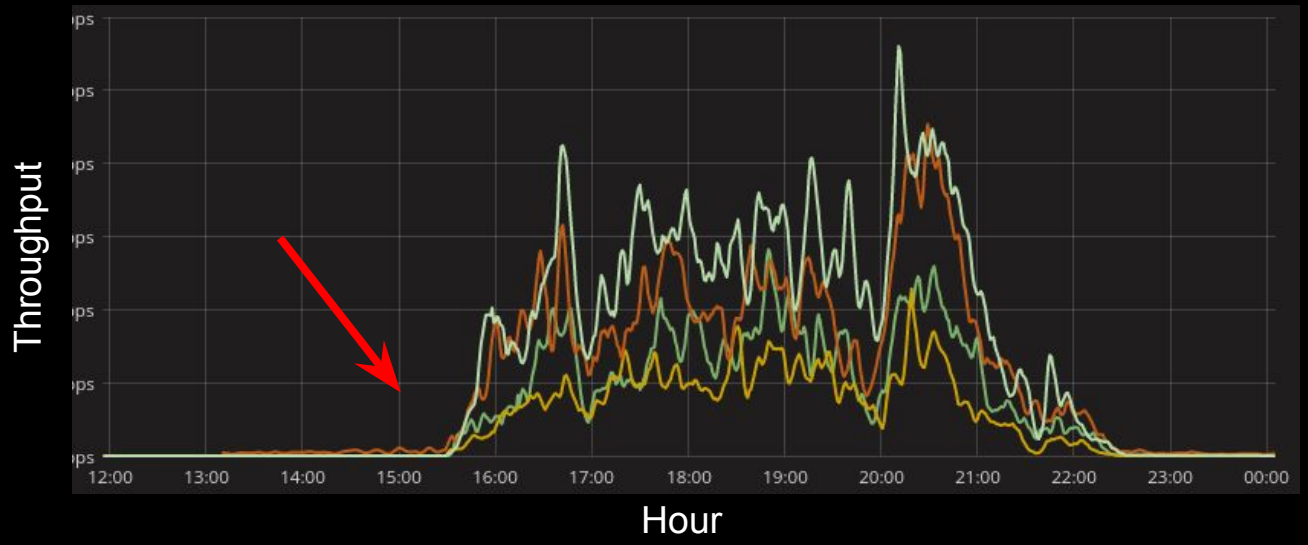

PoP Capacity (%)

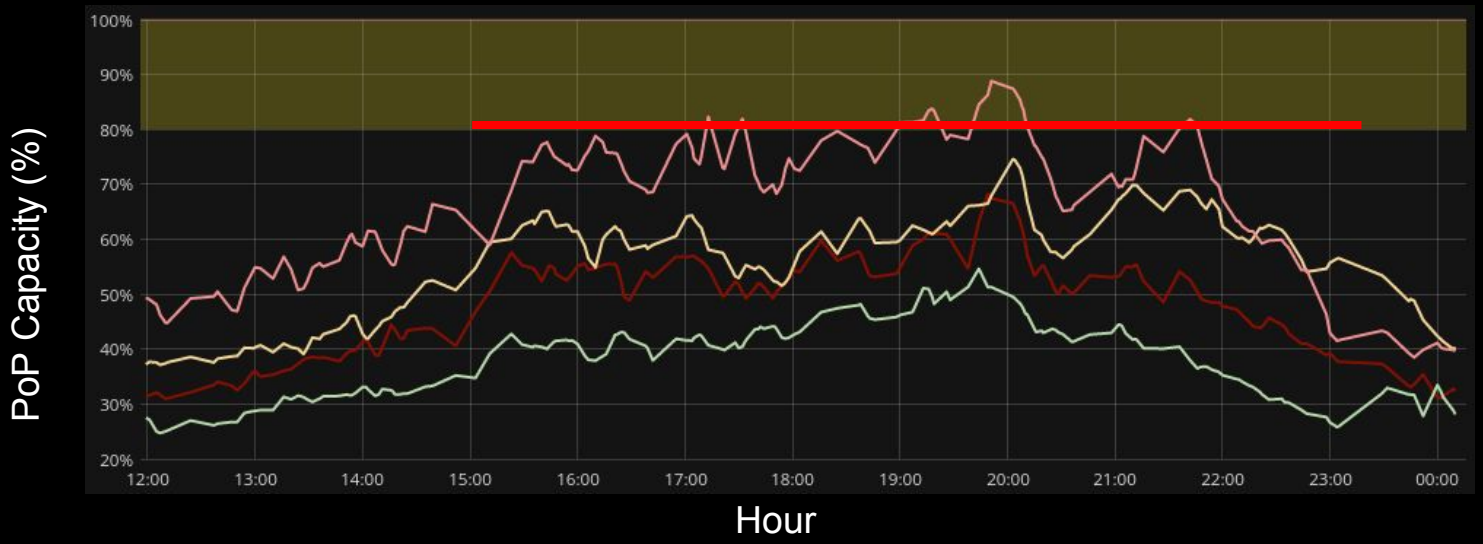

Heteractis in Production

- **Heteractis has been live for almost 2 years.**
  - Making automatic traffic moves nearly daily.
- **Moved from human gated to full auto.**
  - Built confidence of the humans.
  - Widely used as a view into CDN health.
- **Significantly reduced manual human interactions.**

Heteractis

- **We built an automated system for managing traffic.**
- **Implemented in a way that:**
  - Provides visibility into decision making process.
  - Builds trust with humans.

Thank you.